arXiv:1403.6863v1 [cs.LG] 26 Mar 2014

Online Learning of k-CNF Boolean Functions

Joel Veness

DeepMind Technologies Ltd
University of Alberta
joel@deepmind.com

Marcus Hutter

Research School of Computer Science
Australian National University
http://www.hutter1.net/

March 2014

Abstract

This paper revisits the problem of learning a k-CNF Boolean function from examples in the context of online learning under the logarithmic loss. In doing so, we give a Bayesian interpretation to one of Valiant's celebrated PAC learning algorithms, which we then build upon to derive two efficient, online, probabilistic, supervised learning algorithms for predicting the output of an unknown k-CNF Boolean function. We analyze the loss of our methods, and show that the cumulative log-loss can be upper bounded, ignoring logarithmic factors, by a polynomial function of the size of each example.

Keywords

k-CNF, Online Learning, Logarithmic Loss

1 Introduction

In 1984, Leslie Valiant introduced the notion of Probably Approximately Correct (PAC) learnability, and gave three important examples of some non-trivial concept classes that could be PAC learnt given nothing more than a sequence of positive examples drawn from an arbitrary IID distribution [Val84]. One of these examples was the class of k-CNF Boolean functions. Valiant's approach relied on a polynomial time reduction of this problem to that of PAC learning the class of monotone conjunctions. In this paper, we revisit the problem of learning monotone conjunctions from a different viewpoint. This will allow us to derive two new online, probabilistic prediction algorithms that: (i) learn from both positive and negative examples; (ii) avoid making IID assumptions; (iii) suffer low logarithmic loss for arbitrary sequences of examples; (iv) run in polynomial time and space.

Our motivation for investigating this setting comes from our interest in building predictive ensembles for universal source coding. In particular, we are interested in prediction

methods that satisfy the following *power* desiderata, i.e. methods which: (p) make probabilistic predictions; (o) are strongly online; (w) work well in practice; (e) are efficient; (r) and have well understood regret/loss/redundancy properties. Methods satisfying these properties can be combined in a principled fashion using techniques such as those discussed by [VSH12, Mat13], giving rise to ensembles with clearly interpretable predictive capabilities.

Our contribution stems from noticing that Valiant's method can be interpreted as a kind of MAP model selection procedure with respect to a particular family of priors. While this observation is in itself unremarkable, we show that this family of priors possess a number of desirable computational properties. In particular, we show that given n positive examples, it is possible to perform exact Bayesian inference over the 2^d possible hypotheses in time $O(nd)$ and space $O(d)$. Unfortunately, these desirable computational properties do not seem to readily extend to the case where both positive and negative examples are presented. A workable heuristic approach might be to directly use the efficient Bayesian predictor for prediction, only updating the posterior weights when positive instances are seen. However, as well as needlessly throwing away potentially useful information, this approach also makes it impossible to provide meaningful loss guarantees that hold with respect to all input sequences. Our first contribution is to introduce a hybrid algorithm, which uses a combination of Bayesian inference and memorization to construct a polynomial time algorithm whose loss is bounded by $O(d^2)$ for the class of monotone conjunctions. Our second contribution is a more practical algorithm, which requires less space and whose loss is bounded by $O(d \log n)$. Finally, similarly to Valiant, we describe how to combine our algorithms with a reduction that (for fixed k) enables the efficient learning of k-CNF Boolean functions from examples.

2 Preliminaries

We first introduce some notation to formalize our problem setting.

Notation. A Boolean variable x is an element of $\mathcal{B} := \{\bot, \top\} = \{0, 1\}$. We identify false $\bot$ with 0 and true $\top$ with 1, since it allows us to use Boolean functions as likelihood functions for deterministically generated data. We keep the boolean operator notation whenever more suggestive. The unary not operator is denoted by $\neg$, and is defined as $\neg : 0 \mapsto 1; 1 \mapsto 0$ ($\neg x = 1 - x$). The binary conjunction and disjunction operators are denoted by $\wedge$ and $\vee$ respectively, and are given by the maps $\wedge : (1, 1) \mapsto 1;$ or 0 otherwise ($x \wedge y = x \cdot y$). and $\vee : (0, 0) \mapsto 0;$ or 1 otherwise ($x \vee y = \max\{x, y\}$). A literal is a Boolean variable x or its negation $\neg x$; a positive literal is a non-negated Boolean variable. A clause is a finite disjunction of literals. A monotone conjunction is a conjunction of zero or more positive literals. For example, $x_1 \wedge x_3 \wedge x_6$ is a monotone conjunction, while $\neg x_1 \wedge x_3$ is not. We adopt the usual convention with conjunctions of defining the zero literal case to be vacuously true. The power set of a set $\mathcal{S}$ is the set of all subsets of $\mathcal{S}$, and will be denoted by $\mathcal{P}(\mathcal{S})$. For convenience, we further define $\mathcal{P}_d := \mathcal{P}(\{1, 2, \ldots, d\})$. We

also use the Iverson bracket notation $[\![P]\!]$, which given a predicate P, evaluates to 1 if P is true and 0 otherwise. With our identification this is optional but useful syntactic sugar. We also use the notation $x_{1:n}$ and $x_{<n}$ to represent the sequences of symbols $x_1 x_2 \dots x_n$ and $x_1 x_2 \dots x_{n-1}$ respectively. Furthermore, base two is assumed for all logarithms in this paper. Finally, we use the notation a^i to index the ith component of a Boolean vector $a \in \mathcal{B}^d$.

Problem Setup. We consider an online, sequential, binary, probabilistic prediction task with side information. At each time step $t \in \mathbb{N}$, a d-dimensional Boolean vector of side information $a_t \equiv (a_t^1, ..., a_t^d) \in \mathcal{B}^d$ is presented to a probabilistic predictor $\rho_t : \mathcal{B}^d \to (\mathcal{B} \to [0,1])$, which outputs a probability distribution over $\mathcal{B}$. A label $x_t \in \mathcal{B}$ is then revealed, with the predictor suffering an instantaneous loss of $\ell_t := -\log \rho_t(x_t; a_t)$, with the cycle continuing ad infinitum. It will also prove convenient to introduce the joint distribution $\rho(x_{1:n}; a_{1:n})$, which lets us express the cumulative loss $\mathcal{L}_n(\rho)$ in the form

$$\mathcal{L}_n(\rho) \;:=\; \sum_{i=1}^{n} \ell_t \;=\; -\log \prod_{t=1}^{n} \rho_t(x_t; a_t) \;=:\; -\log \rho(x_{1:n}; a_{1:n})$$

We later use the above quantity to analyze the theoretical properties of our technique. As is usual with loss or regret based approaches, our goal will be to construct a predictor ρ such that $\mathcal{L}_n(\rho)/n \to 0$ as $n \to \infty$ for an interesting class of probabilistic predictors $\mathcal{M}$. The focus of our attention for the remainder of this paper will be on the class of monotone conjunctions.

Brute force Bayesian learning. Consider the monotone conjunction $h_{\mathcal{S}}(a_t) := \bigwedge_{i \in \mathcal{S}} a_t^i$ for some $\mathcal{S} \in \mathcal{P}_d$, classifying $a_t \in \mathcal{B}^d$ as $h_{\mathcal{S}}(a_t) \in \mathcal{B}$. This can be extended to the function $h_{\mathcal{S}} : \mathcal{B}^{n \times d} \to \mathcal{B}^n$ that returns the vector $h_{\mathcal{S}}(a_{1:n}) := (h_{\mathcal{S}}(a_1), ..., h_{\mathcal{S}}(a_n))$. One natural Bayesian approach to learning monotone conjunctions would be to place a uniform prior over the set of 2^d possible deterministic predictors that are monotone conjunctions of the d Boolean input variables. This gives the Bayesian mixture model

$$\xi_d(x_{1:n}; a_{1:n}) \;:=\; \sum_{\mathcal{S} \in \mathcal{P}_d} \frac{1}{2^d} \nu_{\mathcal{S}}(x_{1:n}; a_{1:n}), \quad \text{where} \quad \nu_{\mathcal{S}}(x_{1:n}; a_{1:n}) \;:=\; [\![h_{\mathcal{S}}(a_{1:n}) = x_{1:n}]\!] \tag{1}$$

is the deterministic distribution corresponding to $h_{\mathcal{S}}$. Note that when $\mathcal{S} = \{\}$, the conjunction $\bigwedge_{i \in \mathcal{S}} a_t^i$ is vacuously true. From here onwards, we will say hypothesis $h_{\mathcal{S}}$ generates $x_{1:n}$ if $h_{\mathcal{S}}(a_{1:n}) = x_{1:n}$. For sequential prediction, the predictive probability $\xi_d(x_t | x_{<t}; a_{1:t})$ can be obtained by computing the ratio of the marginals, that is $\xi_d(x_t | x_{<t}; a_{1:t}) = \xi_d(x_{1:t}; a_{1:t}) \;/\; \xi_d(x_{<t}; a_{<t})$. Note that this form of the predictive distribution is equivalent to using Bayes rule to explicitly compute the posterior weight for each $\mathcal{S}$, and then taking a convex combination of the instantaneous predictions made by each hypothesis.

The loss of this approach for an arbitrary sequence of data generated by some $h_{\mathcal{S}^*}$ for $\mathcal{S}^* \in \mathcal{P}_d$, can be upper bounded by

$$\begin{aligned}\mathcal{L}_n(\xi_d) \;:=\; -\log \xi_d(x_{1:n}; a_{1:n}) \;&=\; -\log \sum_{\mathcal{S}\in\mathcal{P}_d} \tfrac{1}{2^d} [\![h_{\mathcal{S}}(a_{1:n}) = x_{1:n}]\!] \\ &\leq\; -\log \tfrac{1}{2^d} [\![h_{\mathcal{S}^*}(a_{1:n}) = x_{1:n}]\!] \;=\; d.\end{aligned}$$

Of course the downside with this approach is that a naive computation of Equation 1 takes time $O(n\,2^d)$. Indeed one can show that no polynomial-time algorithm in d for ξ_d exists (assuming P≠NP):

Theorem 1 (ξ_d is #P-complete)**.** *Computing the function $f : \{0,1\}^{n\times d} \to \{0,...,2^d\}$ defined as $f(a_{1:n}) := 2^d \xi_d(0_{1:n}; a_{1:n})$ is #P-complete.*

We prove hardness by a two-step reduction: counting independent sets, known to be #P-hard, to computing the cardinality of a union of power sets to computing ξ_d:

Definition 2 (UPOW)**.** *Given a list of n subsets $\mathcal{S}_1, \ldots, \mathcal{S}_n$ of $\{1, \ldots, d\}$, compute $A := |\mathcal{P}(\mathcal{S}_1) \cup \cdots \cup \mathcal{P}(\mathcal{S}_n)|$, i.e. the size of the union of the power sets of $\mathcal{S}_1, \ldots, \mathcal{S}_n$.*

Lemma 3 (UPOW→ ξ_d)**.** *If a_t is defined as the d-dimensional characteristic bit vector describing the elements in $\mathcal{S}_t$, i.e. $a_t^i := [\![i \in \mathcal{S}_t]\!]$, then $A = 2^d[1 - \xi_d(0_{1:n}|a_{1:n})]$.*

Proof. Since $h_{\mathcal{S}}(a_t)=1$ iff $\mathcal{S}\subseteq\mathcal{S}_t$ iff $\mathcal{S}\in\mathcal{P}(\mathcal{S}_t)$ we have

$$[\![h_{\mathcal{S}}(a_{1:n}) = 0_{1:n}]\!] \iff \bigwedge_{t=1}^{n} [h_{\mathcal{S}}(a_t) = 0] \iff \neg\exists t : \mathcal{S} \in \mathcal{P}(\mathcal{S}_t) \iff \mathcal{S} \notin \mathcal{P}(\mathcal{S}_1)\cup...\cup\mathcal{P}(\mathcal{S}_n)$$

which implies $\sum_{\mathcal{S}\in\mathcal{P}_d} \nu_{\mathcal{S}}(0_{1:n}|a_{1:n}) = 2^d - A$. □

The intuition behind Lemma 3 is that since ξ_d uses a uniform prior over $\mathcal{P}_d$, the number of hypotheses consistent with the data is equal to $2^d\xi_d(0_{1:n}|a_{1:n})$, and therefore the number of hypotheses inconsistent with the data is equal to $2^d[1 - \xi_d(0_{1:n}|a_{1:n})]$. One can easily verify that the set of hypotheses inconsistent with a single negative example is $I_t := \mathcal{P}\left(\{i \in \{1,\ldots,d\} : [\![a_t^i = 1]\!]\}\right)$, hence the set of hypotheses inconsistent with the data is equal to $|\cup_{t=1}^n I_t|$.

Theorem 4 (IS→UPOW, Brendan McKay, private communication)**.** *UPOW is #P-hard.*

Proof. Let $G = (V, E)$ be an undirected graph with vertices $V = \{1, ..., d\}$ and edges $E = \{e_1, ..., e_n\}$, where edges are $e = \{v, w\}$ with $v, w \in V$ and $v \neq w$. An independent set I is a set of vertices no two of which are connected by an edge. The set of independent sets is IS $:= \{I \subseteq V : \forall e \in E : e \not\subseteq I\}$. It is known that counting independent sets, i.e. computing |IS| is #P-hard [Vad01].

We now reduce IS to UPOW: Define $\mathcal{S}_t := V \setminus e_t$ for $t \in \{1, ..., n\}$ and consider any $W \subseteq V$ and its complement $\overline{W} = V \setminus W$. Then

$$\begin{aligned}&W \notin \text{IS} \iff \exists e \in E : e \subseteq W \iff \exists t : e_t \subseteq W \iff \exists t : \overline{W} \subseteq \mathcal{S}_t \\ &\iff \exists t : \overline{W} \in \mathcal{P}(\mathcal{S}_t) \iff \overline{W} \in \mathcal{P}(\mathcal{S}_1) \cup ... \cup \mathcal{P}(\mathcal{S}_n)\end{aligned}$$

Since set-complement is a bijection and there are $2^{|V|}$ possible W, this implies

$$|\text{IS}| \ + \ |\mathcal{P}(\mathcal{S}_1) \cup ... \cup \mathcal{P}(\mathcal{S}_n)| \ = \ 2^{|V|}$$

Hence an efficient algorithm for computing $|\mathcal{P}(\mathcal{S}_1) \cup ... \cup \mathcal{P}(\mathcal{S}_n)|$ would imply the existence of an efficient algorithm for computing |IS|. □

Proof. **of Theorem 1.** Lemma 3 and Theorem 4 show that f is #P-hard. What remains to be shown is that f is in $\#P$. First consider UPOW function $u : \mathcal{P}_d^n \to \{0, ..., 2^d\}$ defined as $u(\mathcal{S}_1, ..., \mathcal{S}_d) := A$. With identification $\{0,1\}^d \cong \mathcal{P}_d$ via $a_t^i = [\![i \in \mathcal{S}_t]\!]$ and $\mathcal{S}_t = \{i : a_t^i = 1\}$, Lemma 3 shows that $f(a_{1:n}) + u(\mathcal{S}_1, ..., \mathcal{S}_n) = 2^d$. Since $\mathcal{S} \in \mathcal{P}(\mathcal{S}_1) \cup ... \cup \mathcal{P}(\mathcal{S}_n)$ iff $\exists t : \mathcal{S} \in \mathcal{P}(\mathcal{S}_t)$ iff $\exists t : \mathcal{S} \subseteq \mathcal{S}_t$, the non-deterministic polynomial time algorithm "Guess $\mathcal{S} \in \mathcal{P}_d$ and accept iff $\exists t : \mathcal{S} \subseteq \mathcal{S}_t$" has exactly A accepting paths, hence u is in $\#P$. Since this algorithm has 2^d paths in total, swapping accepting and non-accepting paths shows that also f is in $\#P$. □

One interesting feature of our reduction was that we only required a sequence of negative examples. As we shall see in Section 3, exact Bayesian inference is tractable if only positive examples are provided. Finally, one can also show that the Bayesian predictor ξ_d obtains the optimal loss.

Proposition 5. *There exists a sequence of side information $a_{1:2^d} \in \mathcal{B}^{2^d \times d}$ such that for any probabilistic predictor $\rho_t : \mathcal{B}^d \to (\mathcal{B} \to [0,1])$, there exists an $\mathcal{S} \in \mathcal{P}_d$ such that $h_{\mathcal{S}}$ would generate a sequence of targets that would give $\mathcal{L}_{2^d}(\rho) \geq d$.*

Proof. Consider the sequence of side information $a_{1:2^d} \in \mathcal{B}^{2^d \times d}$, where a_d^i is defined to be the ith digit of the binary representation of t, for all $1 \leq t \leq 2^d$. As

$$|\{x_{1:2^d} \ : \ x_{1:2^d} \text{ is generated by an } \mathcal{S} \in \mathcal{P}_d\}| = 2^d, \tag{2}$$

to have $\mathcal{L}_{2^d}(\rho) < \infty$ for all $x_{1:2^d}$, we need $\rho(x_{1:2^d}) > 0$ for each of the 2^d possible target strings, which implies that $\mathcal{L}_{2^d}(\rho) \geq d$. □

Memorization. As a further motivating example, it is instructive to compare the exact Bayesian predictor to that of a naive method for learning monotone conjunctions that simply memorizes the training instances, without exploiting the logical structure within the class. To this end, consider the sequential predictor that assigns a probability of

$$m_d(x_n | x_{<n}; a_{1:n}) = \begin{cases} [\![x_n = l(a_{1:n}, x_{<n})]\!] & \text{if } \ a_n \in \{a_1, \ldots, a_{n-1}\}; \\ \frac{1}{2} & \text{otherwise} \end{cases}$$

to each target, where $l(a_{1:n}, x_{<n})$ returns the value of x_t for some $1 \leq t \leq n-1$ such that $a_n = a_t$. Provided the data is generated by some $h_{\mathcal{S}}$ with $\mathcal{S} \in \mathcal{P}_d$, the loss of the above memorization technique is easily seen to be at most 2^d. This follows since an excess loss of 1 bit is suffered whenever a new a_k is seen, and there are at most 2^d distinct inputs (of course no loss is suffered whenever a previously seen a_k is repeated). While

both memorization and the Bayes predictor suffer a constant loss that is independent of the number of training instances, the loss of the memorization technique is exponentially larger as a function of d. This is offset by the fact that memorization can be implemented in essentially $O(nd)$ time by storing the previously seen examples in a hash table. Later we show that by using a combination of Bayesian inference and a slightly more sophisticated form of memorization, it is possible to construct a method whose running time and loss are both bounded by polynomial functions of n and d.

3 Bayesian learning of monotone conjunctions from positive examples

We now show how exact Bayesian inference over the class of monotone conjunctions can be performed efficiently, provided learning only occurs from positive examples $x_{1:n} \equiv 1_{1:n}$. Using the generalized distributive law [AM00] we derive an alternative form of Equation 1 that can be straightforwardly computed in time $O(nd)$.

Proposition 6. *For all $n, d \in \mathbb{N}$, for all $a_{1:n} \in \mathcal{B}^{n \times d}$, then*

$$\xi_d(1_{1:n}; a_{1:n}) \;=\; \prod_{i=1}^{d} \left(\frac{1}{2} + \frac{1}{2} \left[\!\!\left[\bigwedge_{t=1}^{n} a_t^i \right]\!\!\right] \right).$$

Proof. Consider what happens when the expression $\prod_{i=1}^{d} \left(\frac{1}{2} + \frac{1}{2} [\![\bigwedge_{t=1}^{n} a_t^i]\!] \right)$ is expanded. We get a sum containing 2^d terms, that can be rewritten as

$$\sum_{\mathcal{S} \in \mathcal{P}_d} \frac{1}{2^d} \left[\!\!\left[\bigwedge_{i \in \mathcal{S}} \bigwedge_{t=1}^{n} a_t^i \right]\!\!\right] \;=\; \sum_{\mathcal{S} \in \mathcal{P}_d} \frac{1}{2^d} [\![h_{\mathcal{S}}(a_{1:n}) = 1_{1:n}]\!] \;=\; \xi_d(1_{1:n}; a_{1:n}).$$

where the second equality follows from Equation 1 and the first one from

$$\nu_{\mathcal{S}}(1_{1:n} | a_{1:n}) \;=\; [\![h_{\mathcal{S}}(a_{1:n}) = 1_{1:n}]\!] \;=\; \bigwedge_{t=1}^{n} h_{\mathcal{S}}(a_t) \;=\; \bigwedge_{t=1}^{n} \bigwedge_{i \in \mathcal{S}} a_t^i \;=\; \bigwedge_{i \in \mathcal{S}} \bigwedge_{t=1}^{n} a_t^i \qquad (3)$$

□

On MAP model selection from positive examples. If we further parametrize the right hand side of Proposition 6 by introducing a hyper-parameter $\alpha \in (0, 1)$ to give

$$\xi_d^{\alpha}(1_{1:n}; a_{1:n}) \;:=\; \prod_{i=1}^{d} \left((1 - \alpha) + \alpha \left[\!\!\left[\bigwedge_{t=1}^{n} a_t^i \right]\!\!\right] \right), \qquad (4)$$

we get a family of tractable Bayesian algorithms for learning monotone conjunctions from positive examples. The α parameter controls the bias toward smaller or larger formulas; smaller formulas are favored if $\alpha < \frac{1}{2}$, while larger formulas are favored if

$\alpha > \frac{1}{2}$, with the expected formula length being αd. If we denote the prior over $\mathcal{S}$ by $w_\alpha(\mathcal{S}) := \alpha^{|\mathcal{S}|}(1-\alpha)^{d-|\mathcal{S}|}$, we get the mixture

$$\xi_d^\alpha(x_{1:n}; a_{1:n}) \;=\; \sum_{\mathcal{S}\in\mathcal{P}_d} w_\alpha(\mathcal{S})\nu_{\mathcal{S}}(x_{1:n}; a_{1:n})$$

From this we can directly read off the maximum a posteriori (MAP) model

$$\mathcal{S}'_n \;:=\; \arg\max_{\mathcal{S}\in\mathcal{P}_d} w_\alpha(\mathcal{S}|x_{1:n}; a_{1:n}) \;=\; \arg\max_{\mathcal{S}\in\mathcal{P}_d} w_\alpha(\mathcal{S})\nu_{\mathcal{S}}(x_{1:n}|a_{1:n})$$

under various choices of α. The second equality follows from Bayes rule. For positive examples (i.e. $x_{1:n} = 1_{1:n}$), Equation 3 allows us to rewrite this as

$$\mathcal{S}'_n \;=\; \arg\max_{\mathcal{S}\in\mathcal{P}_d} w_\alpha(\mathcal{S}) \left[\!\!\left[\bigwedge_{i\in\mathcal{S}} \bigwedge_{t=1}^{n} a_t^i \right]\!\!\right]$$

For $\alpha > \frac{1}{2}$, the MAP model $\mathcal{S}'_n$ at time n is unique, and is given by

$$\mathcal{S}'_n \;=\; \left\{ i \in \{1,\ldots,d\} : \bigwedge_{t=1}^{n} a_t^i \right\}.$$

For $\alpha = \frac{1}{2}$, a MAP model is any subset of $\mathcal{S}'_n$. For $\alpha < \frac{1}{2}$, the MAP model is $\{\}$. Finally, we remark that the above results allow for a Bayesian interpretation of Valiant's algorithm for PAC learning monotone conjunctions. His method, described in Section 5 of [Val84], after seeing n positive examples, outputs the concept $\bigwedge_{i\in\mathcal{S}'_n} x_i$; in other words, his method can be interpreted as doing MAP model selection using a prior belonging to the above family when $\alpha > \frac{1}{2}$.

A Heuristic Predictor. Next we discuss a heuristic prediction method that incorporates Proposition 6 to efficiently perform Bayesian learning on only the positive examples. Consider the probabilistic predictor ξ_d^+ defined by

$$\xi_d^+(x_n|x_{<n}; a_{1:n}) \;:=\; \frac{\xi_d(x^+_{<n}x_n; a^+_{<n}a_n)}{\xi_d(x^+_{<n}; a^+_{<n})}, \tag{5}$$

where we denote by $a^+_{<n}$ the subsequence of $a_{<n}$ formed by deleting the a_k where $x_k = 0$, for $1 \le k \le n-1$. Similarly, $x^+_{<n}$ denotes to the subsequence formed from $x_{<n}$ by deleting the x_k where $x_k = 0$. Note that since $\xi_d(x^+_{<n}0; a^+_{<n}a_n) = \xi_d(x^+_{<n}; a^+_{<n})\,(1 - \xi_d(1|x^+_{<n}; a^+_{<n}a_n))$, Equation 4 can be used to efficiently compute Equation 5. To further save computation, the values of the $\bigwedge_{t=1}^{n} a_t^i$ terms can be incrementally maintained using $O(d)$ space. Using these techniques, each prediction can be made in $O(d)$ time.

Of course the main limitation with this approach is that it ignores all of the information contained within the negative instances. It is easy to see that this has disastrous implications for the loss. For example, consider what happens if a sequence of n identical negative instances are supplied. Since no learning will ever occur, a positive constant loss will be suffered at every timestep, leading to a loss that grows linearly in n. This suggests that some form of memorization of negative examples is necessary; we will explore this further in the next section.

Discussion. There are some other noteworthy examples of where it is possible to efficiently perform exact Bayesian inference over large discrete spaces. The Context Tree Weighting algorithm [WST95] performs exact model averaging over the space of bounded depth Prediction Suffix Trees [BEYY04]. The switch distribution [vEGdR07] performs exact model averaging over sequences of model indices, where each model index sequence describes a model formed by composing a particular sequence of indexed base models. The Context Tree Switching algorithm [VNHB12] combines these ideas to perform model averaging over all possible sequences of bounded depth Prediction Suffix Trees. [KAW12] showed how to efficiently use model averaging to extend the ideas behind the switch distribution to robustly handle the case where only a small subset of models from a large pool are expected to make good probabilistic predictions. Partition Tree Weighting [VWBG13], Live and Die Coding [WK97] and related methods for piecewise stationary sources [Wil96, GLL11, SM99] also work by efficiently performing model averaging over large discrete spaces defined by transition diagrams [Wil96]. The common theme amongst these techniques is the careful design of priors that allow for the application of either the generalized distributive law [AM00] and/or dynamic programming to avoid the combinatorial explosion caused by naively averaging over a large number of models.

4 An efficient, low loss algorithm for learning monotone conjunctions

In this section we extend the ideas from the previous sections to construct an efficient online algorithm whose loss is bounded by $O(d^2)$. The main idea is to extend the heuristic predictor so that it simultaneously memorizes negative instances while also favoring predictions of 0 in cases where the Bayesian learning component of the model is unsure. The intuition is that by adding memory, there can be at most 2^d times where a positive loss is suffered. Moving the α parameter closer towards 1 causes the Bayesian component to more heavily weigh the predictions of the longer Boolean expressions consistent with the data, which has the effect of biasing the predictions more towards 0 when the model is unsure. Although this causes the loss suffered on positive instances to increase, we can show that this effect is relatively minor. Our main contribution is to show that by setting $\alpha = 2^{-d/2^d}$, the loss suffered on both positive and negative instances is balanced in the sense that the loss can now be upper bounded by $O(d^2)$. We proceed by first describing the algorithm, before moving on to analyze its loss.

Algorithm. The algorithm works very similarly to the previously defined heuristic predictor, with the following two modifications: firstly, the set of all negative instances is incrementally maintained within a set $\mathcal{A}$, with 0 being predicted deterministically if the current negative instance has been seen before; secondly, the ξ_d terms in Equation 5 are

Algorithm 1 $\zeta_d(x_{1:n}; a_{1:n})$

1: $w_i \leftarrow 1$ for $1 \leq i \leq d$
2: $\mathcal{A} \leftarrow \{\}$
3: $\alpha \leftarrow 2^{-d/2^d}$
4: $r \leftarrow 1$
5: **for** $t = 1$ to n **do**
6: Observe a_t
7: **if** $a_t \in \mathcal{A}$ **then**
8: $p_t(1; a_t) \leftarrow 0$
9: $p_t(0; a_t) \leftarrow 1$
10: **else**
11: $p_t(1; a_t) \leftarrow \prod_{i=1}^{d} \frac{(1-\alpha)+\alpha w_i a_t^i}{(1-\alpha)+\alpha w_i}$
12: $p_t(0; a_t) \leftarrow 1 - p_t(1; a_t)$
13: **end if**
14: Observe x_t and suffer a loss of $-\log p_t(x_t; a_t)$.
15: **if** $x_t = 1$ **then**
16: **for** $i = 1$ to d **do**
17: $w_i \leftarrow w_i \; a_t^i$
18: **end for**
19: **else**
20: $\mathcal{A} \leftarrow \mathcal{A} \cup \{a_t\}$
21: **end if**
22: $r \leftarrow p_t(x_t; a_t) r$
23: **end for**
24: **return** r

replaced with ξ_d^α, with $\alpha = 2^{-d/2^d}$. More formally,

$$\zeta_d(x_t | x_{<t}; a_{1:t}) \;\; := \;\; \begin{cases} 1 - x_t & \text{if } a_t \in \mathcal{A} \\ \frac{\xi_d^\alpha(x_{<t}^+ x_t; a_{<t}^+ a_t)}{\xi_d^\alpha(x_{<t}^+; a_{<t}^+)} & \text{otherwise} \end{cases} \tag{6}$$

Complete pseudocode is given in Algorithm 1. The algorithm begins by initializing the weights and the set of negative instances $\mathcal{A}$. Next, at each time step t, a distribution $p_t(\cdot; a_t)$ over $\{0, 1\}$ is computed. If a_t has previously been seen as a negative example, the algorithm predicts 0 deterministically. Otherwise it makes its prediction using the previously defined Bayesian predictor (with $\alpha = 2^{-d/2^d}$) that is trained from only positive examples. The justification for Line 11 is as follows: First note that w_i is always equal to the conjunction of the ith component of the inputs corresponding to the positive examples

occurring before time t, formally

$$w_i \;=\; \bigwedge_{\tau=1:a_\tau\notin\mathcal{A}}^{t-1} a_\tau \quad \text{which by Equation 4 implies} \quad \xi_d^\alpha(x^+_{<t};a^+_{<t}) \;=\; \prod_{i=1}^{d}[(1-\alpha)+\alpha w_i]$$

Similarly $\xi_d^\alpha(x^+_{<t}1;a^+_{<t}a_t)=\prod_{i=1}^d[(1-\alpha)+\alpha w_i a_t^i]$, which by Equation 6 for $a_t\notin\mathcal{A}$ implies

$$\zeta_d(x_t=1|x_{<t};a_{1:t}) \;=\; \frac{\prod_{i=1}^d[(1-\alpha)+\alpha w_i a_t^i]}{\prod_{i=1}^d[(1-\alpha)+\alpha w_i]} \;=\; p_t(1;a_t)$$

Trivially $p_t(x_t;a_t)=1-x_t=\zeta_d(x_t|x_{<t};a_{1:t})$ for $a_t\in\mathcal{A}$ from Line 8. After the label is revealed and a loss is suffered, the algorithm either updates $\mathcal{A}$ to remember the negative instance or updates its weights w_i, with the cycle continuing. The algorithm returns $r=\prod_{t=1}^n p_t(x_1;a_t)=\zeta_d(x_{1:n};a_{1:n})$ as claimed. It requires $O(nd)$ space and processes each example in $O(d)$ time.

Analysis. We now analyze the cumulative log-loss when using ζ_d in place of an arbitrary monotone conjunction corresponding to some $\mathcal{S}^*\in\mathcal{P}_d$. We begin by first proving a lemma, before using it to upper bound the loss of Algorithm 1.

Lemma 7. *For all $d\in\mathbb{N}\setminus\{1\}$, we have that $-\log\left(1-2^{-d/2^d}\right)\le d$.*

Proof. We have that

$$-\ln(1-\mathrm{e}^{-d/\mathrm{e}^d}) \;\le\; -\ln\left(1-\frac{1}{1+d/\mathrm{e}^d}\right) \;=\; \ln\frac{1+d/\mathrm{e}^d}{d/\mathrm{e}^d} \;=\; d+\ln\left(\frac{1}{d}+\frac{1}{\mathrm{e}^d}\right) \;\le\; d. \quad (7)$$

The first bound follows from $\mathrm{e}^{-x}\le\frac{1}{1+x}$. The equalities are simple algebra. The last bound follows from $\frac{1}{d}+\frac{1}{\mathrm{e}^d}\le 1$ for $d\ge 2$. (A similar lower bound $-\ln(1-\mathrm{e}^{-d/\mathrm{e}^d})\ge -\ln(1-(1-d/\mathrm{e}^d))=d-\ln d$ shows that the bound is rather tight for large d). Substituting $d\rightsquigarrow d\ln 2$ in (7) and dividing by $\ln 2$ proves the lemma. □

Theorem 8. *If $x_{1:n}$ is generated by a hypothesis $h_{\mathcal{S}^*}$ such that $\mathcal{S}^*\in\mathcal{P}_d$ then for all $n\in\mathbb{N}$, for all $d\in\mathbb{N}\setminus\{1\}$, for all $x_{1:n}\in\mathcal{B}^n$, for all $a_{1:n}\in\mathcal{B}^{n\times d}$, we have that $\mathcal{L}_n(\zeta_d)\le 2d^2$.*

Proof. We begin by decomposing the loss into two terms, one for the positive and one for

the negative instances.

$$\begin{aligned}
\mathcal{L}_n(\zeta_d) &= \sum_{t=1}^{n} -\log \zeta_d(x_t | x_{<t}; a_{1:t}) \\
&= \sum_{\substack{t\in[1,n] \\ \text{s.t. } x_t=1}} -\log \zeta_d(x_t = 1 | x_{<t}; a_{1:t}) \quad + \sum_{\substack{t\in[1,n] \\ \text{s.t. } x_t=0}} -\log \zeta_d(x_t = 0 | x_{<t}; a_{1:t}) \\
&= \sum_{\substack{t\in[1,n] \\ \text{s.t. } x_t=1}} -\log \frac{\xi_d^\alpha(x^+_{1:t}; a^+_{1:t})}{\xi_d^\alpha(x^+_{<t}; a^+_{<t})} \quad + \sum_{\substack{t\in[1,n] \\ \text{s.t. } x_t=0}} -\log \zeta_d(x_t = 0 | x_{<t}; a_{1:t}) \\
&= -\log \xi_d^\alpha(x^+_{1:n}; a^+_{1:n}) \quad + \sum_{\substack{t\in[1,n] \\ \text{s.t. } x_t=0}} -\log \zeta_d(x_t = 0 | x_{<t}; a_{1:t}), \qquad (8)
\end{aligned}$$

where we have used the notation $[1, d] := \{1, 2, \ldots, d\}$. The final step follows since the left summand telescopes. Next we will upper bound the left and right terms in Equation 8 separately.

For $\alpha \in (0.5, 1)$, we have for the left term that

$$\begin{aligned}
-\log \xi_d^\alpha(x^+_{1:n}; a^+_{1:n}) &\leq -\log\left(\alpha^{|\mathcal{S}^*|}(1-\alpha)^{d-|\mathcal{S}^*|}\right) \\
&\leq -d\log(1-\alpha). \qquad (9)
\end{aligned}$$

Now, let $\mathcal{U} := \left\{t \in [1, n] : x_t = 0 \ \wedge \ \bigwedge_{i=1}^{t-1}(a_t \neq a_i)\right\}$ denote the set of time indices where a particular negative instance is seen for the first time and let

$$\mathcal{D}_t := \left\{ i \in [1, d] : \bigwedge_{\tau=1}^{t-1} \left(\neg x_\tau \vee a^i_\tau\right) \right\} \qquad (10)$$

denote the indices of the variables not ruled out from the positive examples occurring before time t. Given these definitions, we have that

$$\begin{aligned}
\xi_d^\alpha(x^+_{<t}; a^+_{<t}) &= \sum_{\mathcal{S}\in\mathcal{P}_d} \alpha^{|\mathcal{S}|}(1-\alpha)^{d-|\mathcal{S}|} \left[\!\left[h_\mathcal{S}(a^+_{<t}) = x^+_{<t} \right]\!\right] = \sum_{\mathcal{S}\in\mathcal{P}(\mathcal{D}_t)} \alpha^{|\mathcal{S}|}(1-\alpha)^{d-|\mathcal{S}|} \\
&= (1-\alpha)^{d-|\mathcal{D}_t|} \sum_{\mathcal{S}\in\mathcal{P}(\mathcal{D}_t)} \alpha^{|\mathcal{D}_t|}(1-\alpha)^{|\mathcal{D}_t|-|\mathcal{S}|} = (1-\alpha)^{d-|\mathcal{D}_t|} \qquad (11)
\end{aligned}$$

and similarly for $t \in \mathcal{U}$

$$\begin{aligned}
\xi_d^\alpha(x^+_{<t}0; a^+_{<t}a^+) &= \sum_{\mathcal{S}\in\mathcal{P}_d} \alpha^{|\mathcal{S}|}(1-\alpha)^{d-|\mathcal{S}|} \left[\!\left[h_\mathcal{S}(a^+_{<t}a_t) = x^+_{<t}0 \right]\!\right] \qquad (12) \\
&= \sum_{\mathcal{S}\in\mathcal{P}(\mathcal{D}_t)} \alpha^{|\mathcal{S}|}(1-\alpha)^{d-|\mathcal{S}|} \left[\!\left[h_\mathcal{S}(a_t) = 0 \right]\!\right] \geq \alpha^{|\mathcal{D}_t|}(1-\alpha)^{d-|\mathcal{D}_t|}
\end{aligned}$$

The last inequality follows by dropping all terms in the sum except for the term corresponding the maximally sized conjunction $\bigwedge_{t\in\mathcal{D}_t} x_t$, which must evaluate to 0 given a_t,

since $\mathcal{S}^* \subseteq \mathcal{D}_t$ and $t \in \mathcal{U}$. Using the above, we can now upper bound the right term

$$\sum_{\substack{t\in[1,n]\\ \text{s.t. } x_t=0}} -\log \zeta_d(x_t = 0|x_{<t}; a_{1:t}) \overset{(a)}{=} \sum_{t\in\mathcal{U}} -\log \frac{\xi_d^\alpha(x^+_{<t}0; a^+_{<t}a^+)}{\xi_d^\alpha(x^+_{<t}; a^+_{<t})}$$
$$\overset{(b)}{\leq} \sum_{t\in\mathcal{U}} -\log \alpha^{|\mathcal{D}_t|} \overset{(c)}{\leq} \sum_{t\in\mathcal{U}} -d\log\alpha \overset{(d)}{\leq} -d\,2^d \log\alpha. \tag{13}$$

Step (a) follows from the definition of ζ_d and $\mathcal{U}$ (recall that a positive loss occurs only the first time an input vector is seen). Step (b) follows from Equations 11 and 12. Step (c) follows since $|\mathcal{D}_t| \leq d$ by definition. Step (d) follows since there are at most 2^d distinct Boolean vectors of side information.

Now, by picking $\alpha = 2^{-d/2^d}$, we have from Equation 9 and Lemma 7 that

$$-\log \xi_d^\alpha(x^+_{1:n}; a^+_{1:n}) \leq d^2.$$

Similarly, from Equation 13 we have that

$$\sum_{\substack{t\in[1,n]\\ \text{s.t. } x_t=0}} -\log \zeta_d(x_t = 0|x_{<t}; a_{1:t}) \leq -d\,2^d \log 2^{-d/2^d} = d^2.$$

Thus by summing our previous two upper bounds, we have that

$$\mathcal{L}_n(\zeta_d) \;=\; -\log \xi_d^\alpha(x^+_{1:n}; a^+_{1:n}) \;+ \sum_{\substack{t\in[1,n]\\ \text{s.t. } x_t=0}} -\log \zeta_d(x_t = 0|x_{<t}; a_{1:t}) \leq 2d^2.$$

□

Discussion. So far we have always made the assumption that the targets are generated by some unknown monotone conjunction. If this assumption does not hold, it is possible to observe a label which will have a probability of 0 under ζ_d, which will cause an infinite loss to be suffered. The next section will present an algorithm which, as well as being more space efficient, avoids this problem.

5 A more practical approach

Although the loss of Algorithm 1 is no more than $2d^2$ (and independent of n), a significant practical drawback is its $O(nd)$ space complexity. We now present an alternative algorithm which reduces the space complexity to $O(d)$, at the small price of increasing the worst case loss to no more than $O(d \log n)$.

Algorithm 2 $\pi_d(x_{1:n}; a_{1:n})$

$\mathcal{D} \leftarrow \{1, 2, \ldots, d\}$
$r \leftarrow 1$
for $t = 1$ to n **do**
    Observe a_t
    if $\prod_{i \in \mathcal{D}} a_t^i = 1$ **then**
        $p_t(1; a_t) := t/(t+1)$
        $p_t(0; a_t) := 1/(t+1)$
    else
        $p_t(1; a_t) := 1/(t+1)$
        $p_t(0; a_t) := t/(t+1)$
    end if
    Observe x_t and suffer a loss of $-\log p_t(x_t; a_t)$.
    if $x_t = 1$ **then**
        $\mathcal{D} \leftarrow \mathcal{D} \setminus \{i \in \{1, \ldots, d\} \,:\, a_t^i = 0\}$
    end if
    $r \leftarrow p_t(x_t; a_t)\, r$
end for
return r

Algorithm. The main intuition for our next algorithm follows from the loss analysis of Algorithm 1. Our proof of Theorem 8 led to a choice of $\alpha = 2^{-d/2^d}$, which essentially causes each probabilistic prediction to be largely determined by the prediction made by the longest conjunction consistent with the already seen positive examples. This observation led us to consider Algorithm 2, which uses a smoothed of this. More formally,

$$\pi_d(x_t | x_{<t}; a_{1:t}) \;:=\; \frac{t}{t+1} \left[\!\!\left[\bigwedge_{i \in \mathcal{D}_t} a_t^i = x_t \right]\!\!\right] + \frac{1}{t+1} \left[\!\!\left[\bigwedge_{i \in \mathcal{D}_t} a_t^i \neq x_t \right]\!\!\right];$$

where $\mathcal{D}_t$ denotes the indices of the variables not ruled out from the positive examples occurring before time n as defined in Equation 10.

Complete pseudocode for implementing this procedure in $O(d)$ time per iteration, using $O(d)$ space, is given in Algorithm 2. The set $\mathcal{D}$ incrementally maintains the set $\mathcal{D}_t$. Compared to Algorithm 1, the key computational advantage of this approach is that it doesn't need to remember the negative instances.

Analysis. We next upper bound the loss of Algorithm 2.

Theorem 9. *If $x_{1:n}$ is generated by a hypothesis $h_{\mathcal{S}^*}$ such that $\mathcal{S}^* \in \mathcal{P}_d$, then for all $n \in \mathbb{N}$, for all $d \in \mathbb{N}$, for all $x_{1:n} \in \mathcal{B}^n$, for all $a_{1:n} \in \mathcal{B}^{n \times d}$, we have that $\mathcal{L}_n(\pi_d) \leq (d+1)\log(n+1)$.*

Proof. As $x_{1:n}$ is generated by some $h_{\mathcal{S}^*}$ where $\mathcal{S}^* \in \mathcal{P}_d$, we have that $\mathcal{L}_n(\pi_d) = -\log \pi_d(x_{1:n}; a_{1:n})$. We break the analysis of this term into two cases. At any time $1 \leq t \leq n$, we have either: Case (i): $\bigwedge_{i \in \mathcal{A}} a_t^i = \top$, which implies $h_{\mathcal{S}}(a_t) = 1$ for all $\mathcal{S} \subseteq \mathcal{A} = \mathcal{D}_t$. As the data is generated by some h_{S^*}, we must have $\mathcal{S}^* \subseteq \mathcal{D}_t$ and therefore $x_t = 1$, so a loss of $-\log \frac{t}{t+1}$ is suffered. Case (ii): $\bigwedge_{i \in \mathcal{A}} a_t^i = \bot$, where one of two situations occur: a) if $x_t = 0$ we suffer a loss of $-\log \frac{t}{t+1}$; otherwise b) we suffer a loss of $-\log(1/(t+1)) = \log(t+1)$ and at least one element in $\mathcal{A}$ gets removed. Notice that as the set $\mathcal{A}$ is initialized with d elements, case b) can only occur at most d times given any sequence of data.

Finally, notice that Case (ii b) contributes at most at d times $\log(n+1)$ to the loss. On the other hand, $\log \frac{t+1}{t}$ is suffered for each t of case (i) and (ii a), which can be upper bounded by $\sum_{t=1}^{n} \log \frac{t+1}{t} = \log(n+1)$. Together they give the desired upper bound $(d+1)\log(n+1)$. □

We also remark that Algorithm 2 could have been defined so that $p_t(1; a_t) = 1$ whenever $\bigwedge_{i \in \mathcal{A}} a_t^i = 1$. The reason we instead predicted 1 with probability $\frac{t}{t+1}$ is that it allows Algorithm 2 to avoid suffering an infinite loss if the data is not generated by some monotone conjunction, without meaningfully affecting the loss analysis.

6 Online learning of k-CNF Boolean functions

Finally, we describe how our techniques can be used to probabilistically predict the output of an unknown k-CNF function. Given a set of d variables $\{x_1, \ldots, x_d\}$, a k-CNF Boolean function is a conjunction of clauses $c_1 \wedge c_2 \wedge \cdots \wedge c_m$, where for $1 \leq y \leq m$, each clause c_y is a disjunction of k literals, with each literal being an element from $\{x_1, \ldots, x_d, \neg x_1, \ldots, \neg x_d\}$. The number of syntactically distinct clauses is therefore $(2d)^k$. We will use the notation $\mathcal{C}_d^k$ to denote the class of k-CNF Boolean formulas that can be formed from d variables.

The task of probabilistically predicting a k-CNF Boolean function of d variables can be reduced to that of probabilistically predicting a monotone conjunction over a larger space of input variables. We can directly use the same reduction as used by [Val84] to show that the class of k-CNF Boolean functions is PAC-learnable. The main idea is to first transform the given side information $a \in \mathcal{B}^d$ into a new Boolean vector $c \in \mathcal{B}^{(2d)^k}$, where each component of c corresponds to the truth value for each distinct k-literal clause formed from the set of input variables $\{a^i\}_{i=1}^d$, and then run either Algorithm 1 or Algorithm 2 on this transformed input. In the case of Algorithm 1, this results in an online algorithm where each iteration takes $O(d^k)$ time; given n examples, the algorithm runs in $O(nd^k)$ time and uses $O(nd^k)$ space. Furthermore, if we denote the above process using either Algorithm 1 or Algorithm 2 as $\textsc{Alg1}_d^k$ or $\textsc{Alg2}_d^k$ respectively, then Theorems 8 and 9 allows us to upper bound the loss of each approach as follows.

Corollary 10. *For all $n \in \mathbb{N}$, for all $k \in \mathbb{N}$, for any sequence of side information $a_{1:n} \in \mathcal{B}^{n \times d}$, if $x_{1:n}$ is generated from a hypothesis $h^* \in \mathcal{C}_d^k$, the loss of $\textsc{Alg1}_d^k$ and $\textsc{Alg2}_d^k$ with respect to h^* satisfies the upper bounds $\mathcal{L}_n(\textsc{Alg1}) \leq 2^{2k+1} d^{2k}$ and $\mathcal{L}_n(\textsc{Alg2}) \leq \left(2^k d^k + 1\right)\log(n+1)$ respectively.*

It is also possible to provide similar reductions that can probabilistically predict Boolean functions formed from conjunctions or disjunction in $O(d)$ time and space: in the case of conjunctions, we expand the side information to also include d additional terms, each of which corresponds to the negation of each a_t^i; disjunctions are handled by using De Morgan's law by first flipping the side information components and x_t, and then using the reduction technique for conjunctions.

7 Experimental results

We additionally performed some simulations to better understand our methods. The source code to replicate all reported results can be found at: `jveness.info/software/kcnf.zip`.

Tightness of Theorem 8 and Theorem 9. We first investigated the tightness of the upper bounds given in Theorem 8 and Theorem 9. The results are summarized in Table 1. For every $d \in \{2^k\}_{k=1}^{8}$, a random $\mathcal{S} \in \mathcal{P}_d$ was generated by first sampling a $\theta \in [0, 1]$ uniformly at random, and then sampling a hypothesis corresponding to $\mathcal{S} \in \mathcal{P}_d$ distributed according to the probability mass function $f(\mathcal{S}) := \theta^{|\mathcal{S}|}(1-\theta)^{d-|\mathcal{S}|}$. Next a Boolean vector of side information was generated by sampling from the uniform distribution over $\mathcal{B}^d$, with each algorithm making a prediction, suffering a loss, and then seeing the correct label, for a total of $n = 8192$ time steps. This whole process was then repeated 1000 times, with the reported empirical loss corresponding to the maximum loss obtained across each of the 1000 repeats, for all d. Table 1 suggests that the upper bound on the loss given in Theorem 8 is reasonably tight, typically only 2-3 times higher than the highest observed empirical loss in our experiments.

	2	4	8	16	32	64	128	256
Algorithm 1	3.5	10.6	44.4	200.462	808.9	3745.8	15555.7	63623.4
Theorem 8	8	32	128	512	2048	8192	32768	131072
Algorithm 2	11.9	19.5	36.8	75.8	97.1	98.4	111.3	116.0
Theorem 9	39	65	117	221	429	845	1677	3341

Table 1: Empirical loss (in bits) of Algorithm 1 and Algorithm 2 versus their theoretical upper bounds for various dimensions (specific dimensions indicated in the column headings).

Comparison of methods. Additionally, we compared Algorithms 1 and 2 to the work of [DC02], which uses an efficient model averaging technique to combine the outputs of 2^d different discriminative Naive Bayesian classifiers in $O(d)$ time per step; we give a brief overview of the technique as applied to our setting in Appendix A. We denote this method by MADNB (Model Averaging Discriminative Naive Bayes). We used a similar setup as before, but also varied the number of timesteps n. The results are shown in

	2	4	8	16	32	64	128	256
$n = 2048$	265.4	263.9	251.5	259.0	288.3	302.4	390.2	466
$n = 4096$	494.8	507.0	475.7	484.5	490.1	489.2	534.5	616.0
$n = 8192$	933.8	956.8	926.2	961.3	921.9	927.4	809.4	1133.43
$n = 16384$	1817.0	1808.4	1810.3	1807.8	1790.2	1791.7	1821.7	1993.6
$n = 32768$	3549.0	3554.5	3528.8	3538.3	3496.8	3562.8	3646.1	3821.8

Table 2: Empirical loss of MADNB (in bits).

Table 2. While it is difficult to draw too many conclusions, it does appear that the loss of MADNB scales linearly in our setting.

Furthermore, we compared MADNB to Algorithm 2 using the k-CNF reduction on two binary classification tasks, each taken from the UCI repository [BL13]. For the first task we used the Mushroom Dataset, and for the second we used King-Rook vs King Pawn chess endgames. These data sets were chosen by filtering the list of available datasets to those that dealt with binary classification tasks, containing only categorical input features, and who had more than 1000 instances. We were also careful to check that the results were qualitatively the same under different permutations of the instances and switching of the class labels. The results are summarized in Tables 3 and 4. Here we see two main effects. First is that the real world performance of MADNB is much better than the worst case results presented earlier. Secondly, we see the benefit of expanding the model class by increasing k; it seems that for both datasets, $k = 3$ is sufficient for achieving superior or comparable predictive accuracy to MADNB.

Method	Accuracy	Correct	Mistakes	Total Log Loss (bits)
MADNB	93.17	7569	555	3158.01
monotone	48.44	3935	4189	48494.4
1-CNF	48.67	3954	4170	48313
2-CNF	93.27	7577	547	6156.91
3-CNF	**98.58**	**8009**	**115**	**844.51**
4-CNF	97.49	7920	204	1645.48

Table 3: Results on the Mushroom Data Set.

Method	Accuracy	Correct	Mistakes	Total Log Loss (bits)
MADNB	**88.36**	**2824**	**372**	**1366.06**
monotone	52.07	1664	1532	15616.3
1-CNF	54.57	1744	1452	14836.1
2-CNF	80.51	2573	623	6311.06
3-CNF	85.58	2735	461	4311.65
4-CNF	77.47	2476	720	6848.97

Table 4: Results on the King-Rook vs King-Pawn Dataset.

8 Discussion

In terms of our methods practical utility, we envision our technique being most useful as component of a larger predictive ensemble. To give a concrete example, consider the statistical data compression setting, where the cumulative log-loss under some probabilistic model directly corresponds to the size of a file encoded using arithmetic encoding [WNC87]. Many strong statistical data compression techniques work by adaptively combining the outputs of many different probabilistic models. For example, the high performance PAQ compressor uses a technique known as geometric mixing [Mat13], to combine the outputs of many different contextual models in a principled fashion. Adding our technique to such a predictive ensemble would give it the property that it could exploit k-CNF structure in places where it exists.

9 Conclusion

This paper has provided two efficient, low-loss algorithms for probabilistically predicting targets generated by some unknown k-CNF Boolean function of d Boolean variables in time (for fixed k) polynomial in d. The construction of Algorithm 1 is technically interesting in the sense that it is a hybrid Bayesian technique, which performs full Bayesian inference only on the positive examples, with a prior carefully chosen so that the loss suffered on negative examples is kept small. This approach may be potentially useful for more generally applying the ideas behind Bayesian inference or exponential weighted averaging in settings where a direct application would be computationally intractable. The more practical Algorithm 2 is less interpretable, but has $O(d)$ space complexity and a per instance time complexity of $O(d)$, while enjoying a loss within a multiplicative $\log n$ factor of the intractable Bayesian predictor using a uniform prior.

Acknowledgements. The authors would like to thank the following people: Brendan McKay for providing the proof of Theorem 4; Kee Siong Ng, for the suggestion many years ago to investigate the class of k-CNF formulas from an online, probabilistic perspective; and finally Julien Cornebise and Marc Bellemare for some helpful comments and discussions.

References

[AM00] S.M. Aji and R.J. McEliece. The generalized distributive law. *Information Theory, IEEE Transactions on*, 46(2):325–343, 2000.

[BEYY04] Ron Begleiter, Ran El-Yaniv, and Golan Yona. On prediction using variable order Markov models. *Journal of Artificial Intelligence Research*, 22:385–421, 2004.

[BL13] K. Bache and M. Lichman. UCI machine learning repository, 2013.

[DC02] Denver Dash and Gregory F. Cooper. Exact model averaging with naive Bayesian classifiers. In *Proceedings of the Nineteenth International Conference on Machine Learning (ICML)*, pages 91–98. Morgan Kaufmann. The, 2002.

[GLL11] A. György, T. Linder, and G. Lugosi. Efficient Tracking of Large Classes of Experts. *IEEE Transactions on Information Theory*, 58(11):6709–6725, 2011.

[KAW12] Wouter M. Koolen, Dmitry Adamskiy, and Manfred K. Warmuth. Putting Bayes to sleep. In *NIPS*, pages 135–143, 2012.

[KT81] R. Krichevsky and V. Trofimov. The performance of universal encoding. *Information Theory, IEEE Transactions on*, 27(2):199–207, 1981.

[Mat13] Christopher Mattern. Linear and Geometric Mixtures - Analysis. In *Proceedings of the 2013 Data Compression Conference*, DCC '13, pages 301–310, Washington, DC, USA, 2013. IEEE Computer Society.

[SM99] Gil I. Shamir and Neri Merhav. Low Complexity Sequential Lossless Coding for Piecewise Stationary Memoryless Sources. *IEEE Transactions on Information Theory*, 45:1498–1519, 1999.

[Vad01] S. P. Vadhan. The complexity of counting in sparse, regular, and planar graphs. *SIAM Journal on Computing*, 31(2):398–427, 2001.

[Val84] L. G. Valiant. A theory of the learnable. *Commun. ACM*, 27(11):1134–1142, November 1984.

[vEGdR07] Tim van Erven, Peter Grünwald, and Steven de Rooij. Catching Up Faster in Bayesian Model Selection and Model Averaging. *Neural Information Processing Systems (NIPS)*, 2007.

[VNHB12] Joel Veness, Kee Siong Ng, Marcus Hutter, and Michael H. Bowling. Context Tree Switching. In *DCC*, pages 327–336, 2012.

[VSH12] Joel Veness, Peter Sunehag, and Marcus Hutter. On Ensemble Techniques for AIXI Approximation. In *AGI*, pages 341–351, 2012.

[VWBG13] J. Veness, M. White, M. Bowling, and A. Gyorgy. Partition Tree Weighting. In *Data Compression Conference (DCC), 2013*, pages 321–330, 2013.

[Wil96] Frans M. J. Willems. Coding for a binary independent piecewise-identically-distributed source. *IEEE Transactions on Information Theory*, 42:2210–2217, 1996.

[WK97] F. Willems and M. Krom. Live-and-die coding for binary piecewise i.i.d. sources. In *Information Theory. 1997. Proceedings., 1997 IEEE International Symposium on*, page 68, jun-4 jul 1997.

[WNC87] Ian H. Witten, Radford M. Neal, and John G. Cleary. Arithmetic coding for data compression. *Commun. ACM*, 30:520–540, June 1987.

[WST95] Frans M.J. Willems, Yuri M. Shtarkov, and Tjalling J. Tjalkens. The Context Tree Weighting Method: Basic Properties. *IEEE Transactions on Information Theory*, 41:653–664, 1995.

A MADNB Description

Here we describe the MADNB method of [DC02] as applied to our setting. At a high level, the algorithm combines the predictions made by many Naive Bayes classifiers. Each classifier is described by a set of feature indices $\mathcal{S} \subseteq \{1, ..., d\}$, with the feature values at time t defined to be $[\![a_t^i]\!]$ for all $i \in \mathcal{S}$. The MADNB technique aggregates the predictions made by the 2^d possible Naive Bayes classifiers corresponding to some subset of features $\mathcal{S} \subseteq \{1, ..., d\}$.

Our presentation will first describe how the parameters can be learnt online using the KT estimator [KT81] for a Naive Bayes classifier containing only binary features, before moving on to describe the MADNB technique.

Parameter learning. Consider a sequence $x_{1:n} \in \{0,1\}^n$ generated by successive Bernoulli trials. If a and b denote the number of zeroes and ones in $x_{1:n}$ respectively, and $\theta \in [0,1]$ denotes the probability of observing a 1 on any given trial, then $\Pr(x_{1:n}|\theta) = \theta^b(1-\theta)^a$. One way to construct a distribution over $x_{1:n}$, in the case where θ is unknown, is to weight over the possible values of θ. The KT-estimator uses the weighting $w(\theta) := \text{Beta}(\frac{1}{2},\frac{1}{2}) = \pi^{-1}\theta^{-1/2}(1-\theta)^{-1/2}$, which assigns the probability $\text{KT}(x_{1:n}) := \int_0^1 \theta^b(1-\theta)^a w(\theta)\, d\theta$. This quantity can be efficiently computed online by maintaining the a and b counts incrementally and using the chain rule, that is, $\Pr(x_{n+1} = 1|x_{1:n}) = 1 - \Pr(x_{n+1} = 0|x_{1:n}) = (b + 1/2)/(n+1)$.

Online Naive Bayes Classifier. Given binary features described by the indices in $\mathcal{S}$, if we define

$$\text{NB}_{\mathcal{S}}(x_{1:t}, a_{1:t}) \;=\; \text{KT}(x_{1:t}) \prod_{i \in \mathcal{S}} \text{KT}_0(a^i_{1:t}|x_{1:t}) \text{KT}_1(a^i_{1:t}|x_{1:t}) \prod_{i \notin \mathcal{S}} \text{KT}(x_{1:t})$$

the probabilistic prediction made by an online Naive Bayes model at time t is given by

$$\text{NB}_{\mathcal{S}}(x_t|x_{<t}; a_{1:t}) \;:=\; \frac{\text{NB}_{\mathcal{S}}(x_{1:t}, a_{1:t})}{\sum_{y \in \mathcal{B}} \text{NB}_{\mathcal{S}}(x_{<t}y, a_{1:t})},$$

where $\text{KT}_j(a^i_{1:t}|x_{1:t}) := \text{KT}\left(f_j(x_{1:t}, a^i_{1:t})\right)$ and $f_j(x_{1:t}, a^i_{1:t})$ returns the subsequence of $a^i_{1:t}$ described by the indices $1 \le k \le t$ satisfying $x_k = j$ (of length a for $j = 0$ and length b for $j = 1$).

MADNB. The uniform Bayesian mixture over all $\mathcal{S} \subseteq \{1, ..., d\}$ is defined as

$$\text{NB}_*(x_{1:t}, a_{1:t}) \;:=\; \sum_{\mathcal{S} \in \mathcal{P}_d} 2^{-d} \text{NB}_{\mathcal{S}}(x_{1:t}, a_{1:t})$$

The MADNB predictor can now be defined, by at time t assigning the probability

$$\text{MADNB}(x_t|x_{<t}; a_{1:t}) \;:=\; \frac{\text{NB}_*(x_{1:t}, a_{1:t})}{\sum_{y \in \mathcal{B}} \text{NB}_*(x_{<t}y, a_{1:t})}. \tag{14}$$

Equation 14 can be computed efficiently, using the identity

$$\text{NB}_*(x_{1:t}, a_{1:t}) \;=\; \text{KT}(x_{1:t}) \prod_{i=1}^{d} \left[\frac{1}{2}\text{KT}(a^i_{1:t}) + \frac{1}{2}\text{KT}_0(a^i_{1:t}|x_{1:t})\text{KT}_1(a^i_{1:t}|x_{1:t}) \right].$$

Note that the MADNB techniques differs from a pure generative (online as well as offline) Bayesian approach, in the sense that

$$\sum_{\mathcal{S}\in\mathcal{P}_d} \frac{1}{2^d} \prod_{t=1}^{n} \text{NB}_{\mathcal{S}}(x_t|x_{<t}; a_{1:t}) \;\neq\; \prod_{t=1}^{n} \text{MADNB}(x_t|x_{<t}; a_{1:t}) \;\neq\; \frac{\text{NB}_*(x_{1:t}, a_{1:t})}{\sum_{a_{1:n}} \text{NB}_*(x_{1:t}, a_{1:t})}$$

but has the advantage of being computationally tractable.

B List of Notation

Symbol	**Explanation**
$\mathcal{B}$	$\{\text{false}, \text{true}\} = \{\bot, \top\} = \{0, 1\}$
$[\![\mathcal{B}]\!]$	is 1 if $\mathcal{B}$ is true, and 0 if $\mathcal{B}$ is false
i	side information vector index $\in \{1, \ldots, d\}$
t	data item index $\in \{1, ..., n\}$
k-CNF	k-Conjunctive Normal Form
d	side information dimensionality
a	side information $\in \{0, 1\}^d$
a^i_t	Boolean side information i at time t
$x_t \in \mathcal{B}$	binary label at time t
$\mathcal{S}$	subset of $\{1, ..., d\}$; indices of positive literals in monotone conjunction
$h_{\mathcal{S}}(a_t)$	monotone conjunction hypothesis $\bigwedge_{i\in\mathcal{S}} a^i_t$
$\mathcal{S}'_n$	the MAP model at time n, where $\mathcal{S}'_n \subseteq \{1, \ldots, d\}$
ξ^{α}_d	Bayesian mixture over monotone conjunctions for positive data
ξ^{+}_d	heuristic Bayesian predictor for arbitrary data that learns only from positive data
ζ_d	Algorithm 1, the hybrid Bayesian+memorization predictor
π_d	Algorithm 2, the memory efficient monotone conjunction predictor
$\nu_{\mathcal{S}}$	deterministic predictor corresponding to monotone conjunction $h_{\mathcal{S}}$
$w(\mathcal{S})$	prior weight of the monotone conjunction defined by $\mathcal{S} \subseteq \{1, ..., d\}$
α	hyper-parameter controlling bias toward smaller/larger formulas
e	base of natural logarithm
log	binary logarithm
ln	natural logarithm
$\mathcal{L}_n(\rho)$	cumulative log-loss of ρ at time n

$\mathcal{C}_d^k$	the class of k-CNF Boolean formulas that can be formed from d variables
$\mathcal{A}$	the incrementally maintained set of negative examples in Algorithm 1
$\mathcal{D}_t$	indices of variables not ruled out from the positive examples occurring before time t
$\mathcal{P}_d$	the powerset of $\{1, \ldots, d\}$
KT	KT estimator
$\text{NB}_\mathcal{S}$	Naive Bayes predictor, using feature set $\mathcal{S}$
MADNB	Model averaging discriminative naive Bayes predictor